

HEC-IWG FILE SYSTEMS AND I/O WORKSHOP

**Harriet Coverston & John Howard
Distinguished Engineers
Sun Microsystems**

August 16, 2005

Sun's Strategy



Communities Create Markets

Sun-Sponsored Storage Research

Community building starts with an open exchange of ideas and people

Some specific examples

- NFS and the vnode layer in BSD
- University of Michigan CITI
 - > Reference NFS implementations
- Carnegie Mellon Parallel Data Laboratory
 - > NASD/OBSD, data management, administration
- University of Minnesota Digital Technology Center
 - > Intelligent storage

Key Challenges

- Scalable metadata
- Intelligent storage
- Self-managed storage systems
- Search and retrieve capabilities
- Data access
- Data sharing



Challenge: Think Trillions

- Scaling metadata
 - > Distributed metadata
 - > For fault tolerance and performance
 - > Concurrency
 - > One thread opening m files
 - > n threads opening 1 file
 - > n threads opening m files
 - > Trillions of files
 - > Multiple sites and geographies

Challenge: Intelligent Storage

- Move data-centric operations to storage devices
- Explore implications of object-based storage
 - > Quality of service
 - > Security
 - > Data integrity
- Automate storage system management
 - > Storage nodes become peers of compute nodes



Challenge: Automating Management

- Management
 - > Self-managing infrastructure
 - > Ease of repair
 - > Performance management
- Access Preservation
 - > Nodes enter/exit without disruption
 - > Fault tolerant operation
 - > Continuous access
- Data Preservation
 - > Disaster recovery
 - > Versioning for checkpoint support



Challenge: Finding Stuff

- Finding any one file among trillions of files
- Finding anything in the petabytes of data
 - > Data mining
- Multi-dimensional addressing
 - > Indexing, dynamic translation



Challenge: Moving Data

- Parallel data access
- Sparse access within large data set
- Data rearrangement algorithms
- Integrating new file access methods with MPI-IO
- Relaxed POSIX



Challenge: Sharing Data

- Need to share data in heterogeneous environments
 - > pNFS standards for data sharing
 - > Global access
 - > Object-based storage for pNFS



HEC-IWG FILE SYSTEMS AND I/O WORKSHOP

harriet.coverston@sun.com
jhh@sun.com